US Fish and Wildlife Service | Making Decisions with Multiple Objectives _Part 1_

[MUSIC PLAYING]

OK, welcome back. This is Module H, Making Decisions With Multiple Objectives.

So, we're still in the realm of evaluating trade-offs, or optimization, but we've decided to spend particular time on the class of decision problems that involve multiple objectives, because these kinds of problems come up over and over again in natural resource management. So we wanted to spend more time showing you some tools for solving multiple objective problems.

So in this module, what we're going to talk about is how the structured trade-offs between objectives improve decisions, and we're going to talk through a couple of key steps. We're going to talk about simplification of the problem, using some very easy techniques that we'll show you and that you'll get to practice. Then we're going to talk about converting different attributes to one scale so that we can solve the trade-off problems, or using weighting by preference or values in order to do quantitative trade-offs between objectives. And we're going to show you a bunch of different examples from our experience about how to use these different tools.

So why are structured methods so important for multiple objective problems for weighing trade-offs? Well, basically, because we're talking about apples and oranges. We're talking about weighing trade-offs between very different things-- money versus species persistence, human use of landscapes versus species conservation. It's hard to think about these things on the same scale because they just don't exist on the same scale. Again, we're talking about trading apples for oranges, and how do we do that?

There's a lot of information in multiple objective problems to process informally. We have multiple objectives, we have multiple actions, and then we'll have multiple different models that make predictions about how each of our actions impacts each of our objectives. And this becomes even more complex when we have multiple parties involved in the decision making process. So, multiple different people are putting different values on the relative importance of the different objectives.

The structure of these multiple objective trade-off methods improves our thinking and our decision making by focusing on the separate parts. Remember, the idea is that we're decomposing decisions into their separate parts. We think about the objectives first, before the actions and before we build the models. And we focus on those fundamental objectives and make sure we're developing methods that are most likely to get us to our fundamental objectives.

It's also about transparency. Especially when we have multiple parties involved in and interested in the outcome of our decision processes, it's important that we have a clear rationale and some transparency about how we weigh trade-offs. That's going to increase the buy-in and the defensibility of our decision for these different stakeholder groups that are interested in the decision outcome.

So, what do we have to do to get to the point where we can use multiple objective trade-off methods? Well, just like in any other decision problem, we have to start with a clear statement of our objectives with measurement attributes for each of the different objectives that we've identified. We have to start with a viable set of alternatives. We have to have models that describe the consequences of different actions on the different objectives that we care about.

And usually, we would display this information in a consequence table. We first talked about consequence tables way back in the Consequence Module. We're going to spend a lot of time looking at consequence tables again in this module.

Usually in consequence tables, what we're displaying is a fixed outcome, so we don't have a way of dealing explicitly with uncertainty. However, it is important that we consider things like sensitivity analysis in order to account for uncertainty that we have. And I'll talk a little bit about that as we go through.

So, once we've done that-- we've identified our objectives, our alternatives, and we've built models, and we've displayed the consequences of different actions in terms of different alternatives in a consequence table-- then we're ready to use the multiple objective trade-off methods. Now, the first step is simplifying the problem as much as possible. And we're going to talk about three methods for doing that. I'll demonstrate them, and then you'll get a chance to practice them. The three methods are dominated alternatives, irrelevant objectives, and even swaps.

Once we've done that-- once we've simplified the problem as much as we can-- then we have to do one of the following things. We have to reduce the multiple objectives that remain after we've simplified to a single objective, so that we can use single objective methods for solving the problem, for doing the optimization. We can do that by combining some objectives in some creative ways, and we'll talk about ways to do that.

Or we could transform some objectives into constraints. So for instance, rather than minimizing the cost of our actions, we might say, as long as we stay within some given budget, we'll be happy. That's a way of transforming objectives into constraints, and it makes solving the problem easier.

We also, if we can't reduce to a single objective problem, we also have available some quantitative trade-off methods where we actually, quantitatively weigh the trade-offs between the different objectives. So what's involved in these methods is we have to first put the consequences on a common scale. And we do that by using a very general scale for expressing each of the different objectives. And we'll talk about how to do that.

And then what we do is we develop a weight for each objective, which is really a part of the statement of our objectives. In fact, it's not just what are our objectives, but how relatively important are our objectives to us? And we'll talk about some ways of developing objective weights. And then, we basically sum across all of those objectives with the weights. And we're going to talk about a technique that we use to do that, the SMART technique, Simple Multi-Attribute Rating Technique.

Now, if we can't reduce to a single objective problem, or we can't use the quantitative trade-off methods, or for some other reasons, we might try a different approach. We might just analyze the problem, and show graphically the trade-offs. We will talk about how we can find the efficient frontier between two different objectives, and that allows for negotiation between different parties directly to deal with the trade-offs. And I'll show you an example of that.

Let's start with an example, and you're going to see this example all through this module. Let's start by talking a little bit about it.

Imagine a situation where you're considering repair of some impoundment. You're considering four different alternatives, and you have five different objectives. You'd like to minimize cost. You'd like to maximize the environmental benefit that you get from your impoundment. You'd like to minimize disturbance to the impoundment. You'd like to minimize silt run off, and you'd like to maximize water retention.

Now, you're going to do this with one of four different alternatives. The status quo, continue doing what you're doing. You could do a minor repair of the damaged impoundment, or you could do a major repair of the damaged impoundment or you could do an entire rebuild of the damaged impoundments. So you see this impoundment repair example visually in this consequence table with the objectives along the left side and the alternatives along the top.

Now, we've developed some model to express the consequences of each of the different alternatives on each of the different objectives. So when we look at this consequence table, it's not readily apparent what the smartest thing to do is. It depends on how much you care about cost versus environmental benefit, how important is disturbance to you, or water retention. So we have to deal with the fact that there's trade-offs. There's probably not going to be one alternative that's just best on every one of the different objectives.

When we have a consequence table like this, the first thing we need to do is simplify the problem as much as we can using some easy techniques. And we're going to talk about those three techniques, and then you're going to get a chance to try them out.

The first of the techniques is looking for dominated alternatives. The idea of the dominated alternative is where you have one alternative where, compared to that alternative, all other alternatives perform the same or better on all objectives. So we'll show you that in a second.

Another technique for simply reducing the problem is an irrelevant objective, looking for irrelevant objectives. The idea of an irrelevant objective is that it's an objective not that isn't important to you, but that doesn't help you distinguish amongst alternatives, where the performance measures or the objective is the same over all of the alternatives. And so, because of that, it just doesn't help you discern amongst the alternatives. And we'll show you an example of that in a second.

And then the last of these methods for simplifying the problem is an even swap. An even swap is the situation where you can adjust the consequences of different alternatives to render them equal for a given objective by combining some objectives in an obvious way. If we can put two objectives on the same scale, we can do an even swap between them, and then by doing that, we can render one of the objectives irrelevant. And we'll talk about, again, an example of how to do that.

We'll go back to our impoundment repair example, and we're going to start by looking for dominated alternatives. What we're looking for is one alternative where there is some other alternative that performs as well or better than the dominated alternative on each of the objectives.

We look here at our consequence table, and we look at the rebuild option. We notice that it is more expensive than the major repair. The environmental benefits are the same. It causes more disturbance then the major repair. The silt runoff is the same. and the water retention is actually lower, and we'd like to maximize water retention.

So rebuild as an alternative, compared to the major repair alternative, is dominated. So we say that the major repair alternative dominates the rebuild alternative. There's no reason why we should ever consider the rebuild alternative, because the major repair alternative performs as well or better on all of our objectives. So this is a dominated alternative dominated by major repair.

Another way to see this graphically is imagine a situation where we were considering cost and environmental impacts. We would want to choose alternatives that were both low on cost and low on environmental impacts. If we considered alternatives that were, say, just as costly, but had more serious environmental impacts, that wouldn't make sense to really consider those alternatives. So those alternatives would be dominated. And those are the yellow alternatives that you see here in this figure.

So we can eliminate the rebuild alternative, because it is dominated by major repair. So that's the first technique we can use for simplifying a multiple objective problem. We've now simplified it from four alternatives to three.

So then, let's talk about the concept of irrelevant objectives. The idea of an irrelevant objective is this is an objective that simply doesn't help us distinguish amongst our alternatives, because it doesn't vary in any meaningful way across our alternatives.

If we look at the water retention objective here, once we've eliminated that rebuild option, we realize that all of the alternatives remaining perform exactly the same on the water retention objective. And so, the water retention objective simply isn't helping us distinguish amongst the alternatives. And because of that, we can eliminate that as an objective.

So now, we've simplified our problem again. We've simplified it so that now we only have four objectives remaining that we have to deal with and make trade-offs. Maybe there's even a bit more that we can do.

The left method that we're going to talk about for easy simplification of these quantitative multiple objective trade-off problems is the even swap method. The idea of the even swab method is that maybe we have one objective that we can express in the same units as some other objective.

Now, it would be important if we were working on a team or with multiple stakeholders that everyone agreed that, yes, those two objectives can be expressed in the same terms. If we can do that, we can set the first objective to the same value for alternatives by converting the difference into that second objective that can be expressed in the same units. That leaves us with the first objective being irrelevant.

Imagine a situation where we can think of buying ourself out of the silt runoff problem. We have differences in the silt runoff across our different alternatives, but what if we could invest some money in such a way to mitigate the silt runoff problem? And in fact, that mitigation costs us a half a million dollars per 1,000 cubic feet of silt.

Well, there's another objective that is already expressed in terms of dollars, which is of course the cost objective. We can imagine taking the difference in the silt runoff across the different alternatives and translating that difference into dollars, and then adding that into the cost objective. So for example, for the status quo objective, we're talking about a difference between 5,000 cubic feet for the status quo and 1,000 cubic feet for the minor repair of runoff.

So we can take that difference of 4,000 cubic feet, and we can translate that into dollars. That's going to cost us $2 million to buy ourselves out of that. We can add that cost into the cost of the status quo objective. We do the same thing with the major repair, and then all of the alternatives on the silt runoff objective have the exact same effect.

So that's now an irrelevant objective. We can now eliminate the silt runoff objective from our problem, but we have to carry that even swap information in the cost objective. So the values for the cost objective will now be different.

We now have a reduced problem. We've been able to reduce it by one alternative and by two objectives. But from here, it's not obvious where to go next. Once we've dealt with these easy ways of simplifying the problem, we're going to talk about some more methods for actually dealing with the trade-offs that remain.

But before we do that, stop the video, and you're going to have a chance to try the dominated alternative, the irrelevant objective, and the even swap methods with an example.

You got a chance to practice the dominated alternative, even swap, and irrelevant objective methods to simplify a fairly big multiple objective decision problem, and this is a rolling thunder decision problem. So let's just walk through the solution together and see if everyone got the same kind of information out of it.

Well, the first thing we notice is that the Fall burn is dominated by the Spring burn. And we notice that, because Fall burn costs the same, has the same rancher revenue, same neighbor complaints, does less well in maintaining cover for birds, does less well on effects on listed plants, does less well on effects on butterflies, and does less well on effects on beetles.

So if Spring burn is a possibility, doing a Fall burn would never be sensible, because you'll always do at least as well or better by doing a Spring burn. So Fall burn is dominated by Spring burn, and we can eliminate Fall burn from consideration. So that's a dominant alternative.

Now, once we've done that, we might also notice that now maintain cover for birds as an objective becomes irrelevant. All of the remaining alternatives perform exactly the same on maintain cover for birds. So we can eliminate that objective.

Now, we need to consider even swaps. And in the problem, it was expressed that rancher revenue is worth $120 per unit, so we can translate rancher revenue into costs. So the benefit that we get in terms of rancher revenue is currently expressed in units, but we can express it in terms of dollars. And in fact, it's worth $6,000-- 50 units multiplied by $120 per unit.

So that revenue actually reduces the cost of the action that we're going to take. So we take the 50 units off of rancher revenue, and we then also eliminate that from cost. So we end up with a true cost of $1,000.

And then, once we've done that, rancher revenue becomes an irrelevant objective, because we've zeroed it out across the different alternatives. So we can eliminate that. But we have to carry that new cost and maintain that revised cost of $1,000 throughout the rest of the problem.

Once we've done that, we notice that the no action alternative is a dominated alternative. It's dominated because it is more expensive than grazing, has the same number of neighbor complaints, the same effect on listed plants, does less well for butterflies, and does less well for beetles. So no action is now dominated by the grazing alternative. We can eliminate the no action alternative.

Using these three methods-- dominated alternatives, irrelevant objectives, and even swaps-- we've been able to eliminate two alternatives and two objectives from consideration. We don't have any more places where we can use these methods to simplify the problem, so now we're going to have to grapple with the trade-offs more directly.

So now, we've used the techniques we talked about-- dominated alternatives, irrelevant objectives, and even swaps-- to simplify the problem as much as we can, but we still often will have to grapple with some remaining trade-offs that we can't simplify any further. We're going to not talk about three methods for solving multiple objective problems after we've simplified them as much as we can.

We're going to talk about, first, reducing multiple objectives to a single objective problem by combining objectives or by transforming some objectives into constraints. We're also going to talk about using quantitative trade-off methods for analyzing trade-offs directly, where basically what we do is we put consequences on a common scale and then weight the different objectives that we have to express the relative importance of those objectives. And then, we're also going to talk about the idea of finding the efficient frontier and graphically displaying that kind of information to facilitate negotiation amongst parties.

Let's start to start talking about reducing multiple objectives into a single objective. We can creatively combine objectives. One way of doing this is to sort of price out. The idea here is it's very much like even swaps. It's a way to express one objective in terms of another.

But typically, when we do this, it's going to be somewhat different than even swaps, because it's going to require a bit more negotiation to decide exactly how to express one objective in terms of another. Even swaps work when there's a very obvious way to express one objective in terms of another, but sometimes we need to be a bit more creative in how we do that.

We can also turn some objectives into constraints. The idea here is that we seek to satisfy rather than to optimize some of our objectives. A classic case would be rather than try to minimize the cost of the management actions that we're doing, we simply seek to stay within some fixed budget, and that makes finding a solution easier.

Let's talk about creatively combining objectives, and we'll use the mallard harvest example. The objectives for mallard harvest in North America are to provide substantial harvest opportunity for hunters interested in harvesting mallards, to conserve the mallard population indefinitely, and also to meet the North American Waterfowl Management Plan goal of the number of mallards in any given year.

We can state those quantitatively as maximize harvest year t, maintain the population size, Nt greater than zero for all years, t, and minimize the difference between the goal and Nt, the population size, in any given year when the population size is below the goal of the North American Waterfowl Management Plan.

Let's look at this in a different way, graphically. The way to do this, to combine these, is to discount the value of harvest when we're below the North American Waterfowl Management Plan goal and to consider that for all years out into the future.

Now, what we've done is we've maximized the cumulative long term harvest. We've made sure that the ducks will be around into the future, because of course, if our duck population goes to zero, our harvest will go to zero. And we've also discounted the value of the harvest when the population size is below the North American Waterfowl Management Plan goal. So you can see in this graphic, what we do is we use a utility function to express the value of the harvest when the population size is below the North American Waterfowl Management Plan goal.

What we do is we multiply the harvest by some number. That number is 1 as long as we're above the North American Waterfowl Management Plan population goal, but it is less than 1 when we're below that goal. So we're discounting the harvest when we're below the population goal.

In this way, we've combined multiple objectives into one objective. We can express it in those terms as the product of that utility and the harvest well into the future. Now, once we've done that, we then have a bunch of solution methods available to us. In particular, what's used is stochastic dynamic programming for optimizing the harvest rate in any given year given the population size.

We might think of another way of doing this with a decision tree. We have some stochasticity, and so we'd like to use a decision tree. But we have a hard time using a decision tree with multiple objectives, of course. So what we do is we combine multiple objectives into one objective.

Imagine a situation where we're interested in whether or not we should establish a reserve. So the decision problem is whether or not is to establish a reserve. We make that decision when the species is extant.

Once we establish the reserve, the population may go extinct or it may persist. So there are four different outcomes. We can establish a reserve, and the population can go extinct or it can persist. And we can not establish a reserve, and the population can go extinct or it can persist.

So let's imagine a situation where we establish a reserve, and the population persists. We'd like that outcome, because the species is persisted, but we also recognize that we've spent some money on building the reserve. And so maybe we give that a score of, say, 70.

Imagine that we didn't establish the reserve, but the population still persisted. Compared to that 70, maybe we'd give that a score of 100 to recognize that we still have the species around, but we didn't have to spend the money to establish the reserve. That's our ultimate outcome.

A really bad outcome might be not establishing a reserve and still having a species go extinct. And that seems like a bad outcome, because from a public relations standpoint, it looks pretty bad to have stood around and done nothing while a species goes extinct. So our interested public is not going to be very happy with that outcome. So we give that maybe a score of zero.

And slightly better than that is the situation where the species goes extinct, but at least we tried, and we did what we could by establishing a reserve. So the different scores that you see on this decision tree are associated with three different objectives-- public relations, cost, and species persistence. And we've combined these in creative ways to come up with one measure.

And the way we would actually do this in practice is we would sit with a manager, and we would say, which of these four outcomes is best to you, and why? And then, we would give that best outcome a score of 100. So in this case, that would be the situation where we don't establish the reserve, and the species persists anyway.

Then we would asked the manager, which of these outcomes is worse for you, and why? And in this case, that would be the situation where we don't establish a reserve, and the species goes extinct. And we would give that a score of zero.

And then we would talk with a manager to see how relatively attractive the other outcomes were compared to that 100 and that zero. And by doing that, we could combine these multiple different objectives. And we would call that a multiple attribute utility.

So those are some examples of ways where we can combine multiple objectives into a single objective using some creative methods. Sometimes we can't or don't want to do that, and instead, we actually have to decide to deal with the trade-offs directly using some quantitative trade-off methods.

There are variety of different methods that are available that basically turn multiple objectives into a single objective problem by assigning weights to each objective, and then calculating the weighted score across all the objectives for each alternative. There are a number of different methods available, but the one that we're going to demonstrate today, and the one that the developers of this course are most frequently using, is a method called SMART, the Simple Multi-Attribute Rating Technique. And we're going to demonstrate SMART now, today, and you're going to have a chance to practice it on your own.

The idea of SMART is that we start by normalizing all of our attributes to a 0 to 1 scale, and then we assign weights to each attribute. And we're going to talk about how we might assign weights. Then we calculate a weighted sum of scores for each alternative, and then the recommendation, the recommended best alternative, is the alternative that has the highest weighted score.

Now, we would also, once we do that, want to do some sensitivity analysis. What if we actually have some uncertainty in the model predictions that we've made, or maybe even in our weights? We might want to then ask the question, after we've developed a recommendation, if the uncertainty were resolved differently, would that change the decision that we make? And that might encourage us to do some further analysis, see if we can reduce that uncertainty, and so on.

We're going to use the impound repair example again to go through the process of using the SMART technique. Remember that we simplified as much as we can using the dominated alternatives, irrelevant objectives, and even swaps, so that we now have three alternatives and three objectives. But now, we have to deal with those trade-offs in using the quantitative trade-off technique.

So you see, then, on the spreadsheet, the first consequence matrix here is the same as what you've seen so far. It has the actual values for the different objectives, for each alternative, in the cells. For example, the cost for the status quo alternative is now $2 million.

Below that here, what we have is we have a normalized version of that. So what we've done basically is the best situation for each objective, the best alternative, gets a score of 1, and the worst gets a score of 0. So for the cost, the most expensive alternative, the major repair, is a 0. And the best, the status quo-- and actually, the minor repair-- are tied on this one. Gets a 1.

And what we do is we use some equations to do that, and the equation is here. If we're trying to maximize a given objective, we use this equation. The normalized score for that objective is equal to the value for that objective, minus the minimum for any alternative in that objective, divided by the maximum for any alternative in that objective. Subtract from that the minimum from any alternative and that objective.

So for example, here it would be C6, this cell here, minus the minimum from C6 to E6, divided by the maximum from C6 to E6, minus the minimum again from C6 to E6. That gives us our score there.

Now, if we're trying to minimize rather than maximize the objective, we just use that same formula but start with 1 minus. And you can see that here. 1 minus that exact same quantity gives us the 1 for the status quo, which is the cheapest, and the 0 for the major repair, which is the most expensive.

Now, once we've done that, we've put everything on the same scale, this normalized scale. But now we recognize-- and that part of the statement of our objectives-- is that some objectives are actually more important to us than others. And we integrate that weight at this point. So now, let's look at the part of the spreadsheet that shows the weighted scores, this part here.

Here, what we've done is we've taken each of the cell-specific scores for this normalized version of the consequence table and multiplied those by a weight for each of the objectives. Well, where do we get the weights for our different objectives? Here, what we've said is that environmental benefits are twice as important to us as cost and disturbance are. We've given that a weight of 0.5, and we've given cost and disturbance a weight to 0.25. So in other words, environmental benefits are twice as important to us.

We're going to talk in a minute about how to get those weights, but for, now let's just put some different weights in and see what the impact is. We multiply the score for the status quo by the weight for each of the different objectives, and that gives us a weighted score for each of the different alternatives under each objective.

Now, it's reasonable for us to take those and sum the weighted scores over each alternative. And we've done that here. The alternative with the highest weighted score is the preferred alternative. And we can see that different weights will actually change the situation.

So for example, if we actually cared almost entirely about cost-- say, let's give that a score of 0.8 and 0.1 for the other two. In fact, in that case, we would actually still prefer the minor repair. So we might do something like this once we've decided on weights. We might try different weights if we have some uncertainty in what the weights actually should be or what our decision makers really care most about. And we can use sensitivity analysis to see whether our decision is robust to the sensitivity that we have in our weights, or even sensitive to the uncertainty that we have in our models, in the actual predictions, that show up in the original consequences table.

But this is the SMART technique. What we've done is we've normalized all of the objectives and put them on a scale from 0 to 1, and then we've weighted the objectives, and then we've summed the weighted scores for each alternative over each objective to develop a single measure of the value of that alternative given the weights that we put on them. And that's what we see here in the spreadsheet.